ECSE-489: TELECOMMUNICATION NETWORK LAB

Experiment 5

Milos Prokic, ID#9925529
Mario Prsa, ID#0023631

Table of Contents

Table of Figures and Tables

Abstract

Nowadays, local area networks (LANs) are becoming part of our every day life. Most of us at work and school depend on their flawless operation. But nothing happens by magic – well thought algorithms and protocols are the driving force behind prompt delivery of user's requests. The two major ones are the Ethernet and the token ring multiple access protocols. Neither of the two are a perfect choice, but they both have their advantages and drawbacks, and we will touch on both of these issues. The path that a single bit takes from a sender to receiver is determined by means of routing algorithms. Putting these algorithms to a test was a major issue in this laboratory session. By changing traffic flows, delays and other parameters we have observed how they behave and recover in order to keep the network intact. Open Short Path First (OSPF) is the algorithm we tested.

Introduction

One of the main interests in local area networks (LAN) is the multiple-access control protocol. It is necessary to determine how the different interconnected station nodes will effectively and simultaneously use the shared network connections. A technique that attains this is referred to as a multiple-access protocol[1],[2]. The present experiment required us to investigate two such protocols; Ethernet's CSMA/CD (carrier sensing/collision avoidance) as well as the token ring.

In CSMA/CD, each node is able to sense the shared cable's carrier level before and as it is transmitting a packet. With all the nodes doing so, they strive to avoid transmitting packets on the network line simultaneously and hence avoid possible collisions[1],[2]. However, due to the fact that the nodes are most often distant from each other on the network line, packet transmissions won't always be sensed by all the nodes at the same time and packet collisions will occur[1]. Accordingly, in the first part of the experiment, we examined how certain network dynamics affect the overall utilization and the delay of packets on an Ethernet LAN by the use of previously made OPNET scenarios.

In addition, identical experiments were conducted for a token ring network that constitutes another popular approach to LAN configurations. In such cases, the multiple-access quandary is resolved by passing a token around the network ring and having a single node transmit at any time; the one in the possession of the token. In the case where a node has no data to transmit it simply passes on the token to subsequent nodes[1],[2]. Such a protocol completely avoids the occurrence of collisions[1]. For the purposes of this experiment our network topology was a star rather than a ring topology and identical experiments were carried out as for the case of the Ethernet CSMA/CD protocol.

As a second part of our experiment, we studied an IP routing protocol; the OSPF (open shortest path first). OSPF is a link state protocol that uses Dijkstra's algorithm to determine the least-cost paths in a network[1]. As such, each router constructs a topological map of an entire autonomous system by gathering information about the "shortest" paths from each to every other router. Furthermore, OSPF has now restrictions on how the link weights are assigned[1]. This experiment hence required us to study two OPNET scenarios where link costs were on the one hand inversely proportional to link capacity and on the other were judiciously assigned by ourselves. We examined the dynamics of this protocol in the context of intra-AS routing between gateway routers of four LANs.

After presenting a detailed methodology describing the performed experiments as well as how we carried them out, we will analyze the collected results as well as provide a relevant thorough discussion. We terminate this report with a brief conclusion summarizing our results and discussing important encountered issues.

Methodology

Part 1: LAN Simulations

Two OPNET scenarios were provided where five station nodes were interconnected via a hub, constituting as such a smaller LAN. The exact topology is illustrated in Figure 1.

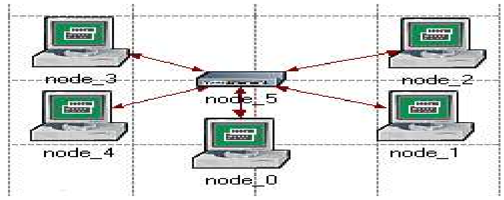

Figure 1: LAN topology used in the OPNET scenario for Part 1

In the first scenario, an Ethernet LAN was implemented and hence using the CSMA/CD multiple access protocol, whereas, a taking-turns protocol; the token-passing was implemented in the second scenario. A predefined traffic of packets was also applied between the five nodes. This default traffic load was fixed by assigning values to the inter-arrival time of packets as well as the packet size. Accordingly, the inter-arrival time was exponentially distributed with mean 0.002 for the token ring and 0.0004 for Ethernet. The packet size was set to exponential~1024 for both scenarios.

The first interest was to examine how increasing the hop propagation delay of all connections in the LAN affects the achieved efficiency (utilization) at every station. This propagation parameter was changed on the links for the Ethernet scenario and at the station nodes for the token ring scenario. The observed parameter was subsequently the global utilization statistic for both networks.

Secondly, for both LAN scenarios, we maintained the propagation delay of all links at a fixed value of 3.3E-6. It was then desired to observe how the average packet delay varies as a function of traffic load. In order to observe and obtain relevant results, we had to vary a traffic parameter and consequently collect values on the average network traffic load and the average network packet delay. As a result, we varied the inter-arrival time of packets parameter at each node of both networks and observed the global delay and received traffic statistics.

Part 2: Routing Experiments

The OPNET modeler was used for the purposes of the second part of the experiment as well. More precisely, we were provided with two new scenarios both reproducing interconnections between gateway routers of four different LANs. In such a constructed autonomous system (AS), an intra-AS routing protocol was studied; the OSPF protocol. The four gateway routers were labeled A, B, C and D and specific IP traffic was predefined for both scenarios as summarized in Table 1. Table 2 identifies furthermore all the present intra-AS bidirectional links with their corresponding data rates.

SCENARIO 1		
Source Node	*Destination Node*	*Traffic Rate*
A	**D**	**42 Mbps**
C	**A**	**35 Mbps**
A	**C**	**22 Mbps**
B	**A**	**20 Mbps**
D	**B**	**20 Mbps**
D	**C**	**20 Mbps**
SCENARIO 2		
Source Node	*Destination Node*	*Traffic Rate*
D	**B**	**23 Mbps**
C	**B**	**23 Mbps**

SCENARIO 1	
Bidirectional link	*Link rate*
A-B	**OC1**
A-C	**DS3**
A-D	**DS3**
B-C	**OC3**
C-D	**OC3**
SCENARIO 2	
Bidirectional link	*Link rate*
A-B	**OC1**
A-C	**DS3**
A-D	**DS3**
B-C	**OC3**
B-D	**DS3**
C-D	**OC3**

Table 1: IP traffic definitions for the OPNET OSPF project *Table 2: Link descriptions for the OPNET OSPF project*

General network routing dynamics had to be examined in order to identify intrinsic characteristics of the OSPF protocol. As such, we used various OPNET functionalities and viewed relevant statistics in order to determine the protocol's convergence time, identify the routes assigned to the IP traffic flows and observe link utilizations. Variations were then made to the network where we increased the data rate of one link, failed another and assigned specific costs to each node interface. Subsequently, by running multiple simulations for the modified scenarios we were able to draw conclusions about OSPF protocol's operation. All completed experiments and collected results are discussed in detail in the subsequent section.

Experiments and Results

Part 1: LAN Simulations

In the simulated Ethernet LAN illustrated in Figure 1, the nodes are interconnected by a broadcast channel where a transmitted frame by a node will be received by all other nodes. In such a high traffic environment, the used multiple access protocol has the responsibility of ensuring that no nodes transmit simultaneously and moreover that there's an efficient use of the network by all nodes. Ethernet's CSMA/CD protocol is hence specifically designed so the nodes make every effort to avoid packet collisions and in addition keep the efficiency at a maximum. By sensing the channel voltage levels, each node can avoid transmitting packets while other ones are doing so and can moreover stop the transmission when it detects a collision[1]. In the case of a collision, the CSMA/CD protocol implements an algorithm directing the behavior of all nodes in order to maintain a high network efficiency.

As an initial experiment, in the provided OPNET scenario we gradually varied the propagation delay of all links in the network and observed the achieved utilization. Figure 2 depicts the latter statistic for various modified scenarios that differ uniquely in the assigned link propagation delay. As can be seen from this plot, as the link propagation delay decreases, the network efficiency or utilization increases. Since the efficiency of Ethernet is defined as fraction of time during which frames are being transmitted without collisions, the observed phenomenon is thus understood by the fact that when the propagation delay decreases, the colliding nodes will abort transmission much faster and waste less the channel. Furthermore, Ethernet's efficiency formula was derived to be:

$$Efficiency = \frac{1}{1 + 5t_{prop}/t_{trans}} \text{ [3]}$$

As we can see, if the propagation delay approaches zero the efficiency will tend towards one. This can also be seen from the plot in Figure 2 where for a propagation delay of 1e-6 the efficiency is indeed very high (75%). Hence, we can conclude that the efficiency of CSMA/CD can easily approach 100% for low propagation delays.

The same experiment was repeated for the token ring LAN; the results are shown in Figure 3. Global statistics regarding utilization however couldn't be collected, therefore the utilization parameter was evaluated for the hub node. The latter can be considered as a judicious choice since all the traffic transits through that node. We can draw the same conclusion as for Ethernet's case where network utilization tends to increase for lower propagation delays. However, the achieved efficiency for identical low propagation delays proved to me much lower for the token ring LAN, as can be seen by comparing figures 2 and 3. This is as expected due to the nature of the multiple access protocol in the token ring, which is a taking-turns protocol. Therefore, at moments nodes can have frames to transmit but are unable to do so since they're waiting for the token to be passed to them, resulting in a decreased efficiency of the network.

In a second experiment using the same OPNET scenarios, it was of interest to compare the average packet delay as a function of traffic load. The procedure followed was as described in the methodology section. To obtain however precise results regarding the average packet delay statistic, ideally our simulation had to run to infinity or at least for a large amount of time in order for the mentioned parameter to converge to a fixed value. We ran however our simulation for ten minutes and obtained acceptable results. The average packet delay as a function of the overall received traffic is plotted in figures 4 and 5 for the Ethernet and token ring scenarios respectively.

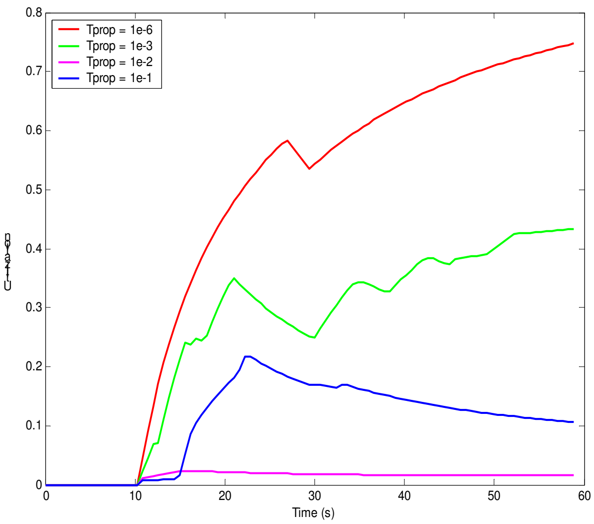

Figure 2: Ethernet network utilization over time for various propagation delay values

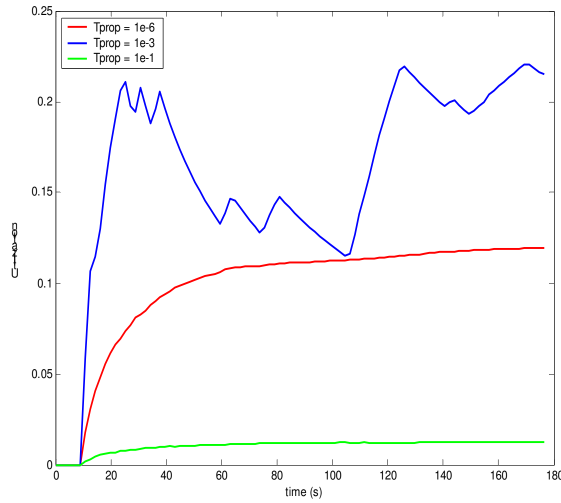

Figure 3: Token ring network utilization over time for various hop propagation delay values

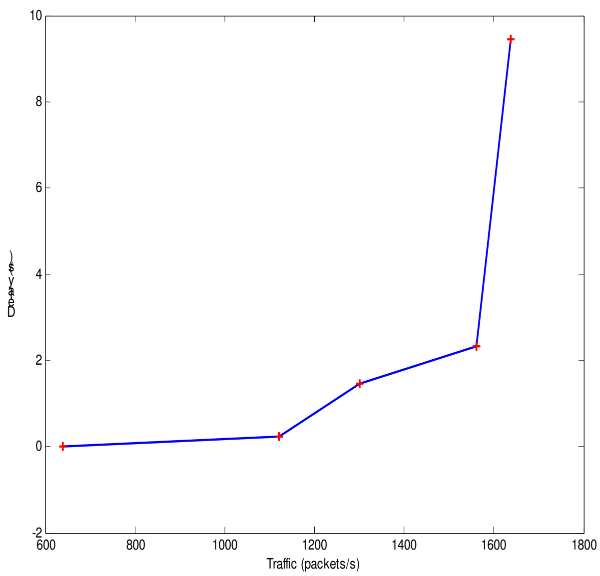

Figure 4: Average packet delay versus received traffic plot for the Ethernet LAN

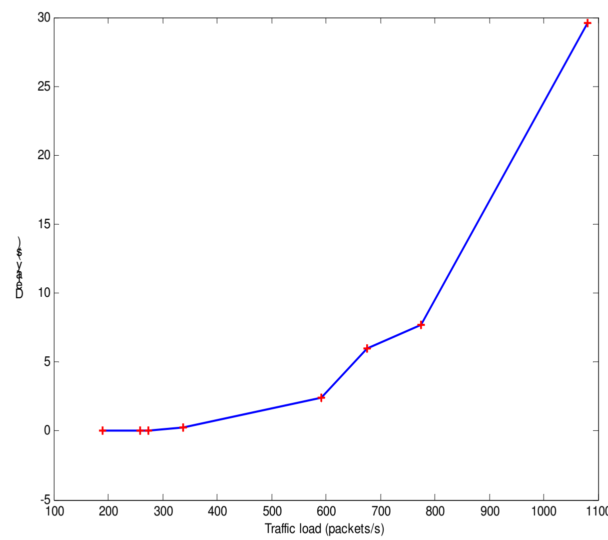

Figure 5: Average packet delay versus received traffic

plot for the token ring LAN

As can be seen, the average packet delay remains relatively low (close to zero) for small traffic loads and then tends to increase exponentially as the traffic load rises above a certain threshold value. In the Ethernet case, this is due to the fact that as the traffic on the network increases, more collisions will occur causing retransmission of frames and consequently increasing their overall delay. The traffic threshold value can be evaluated to be roughly 1500 packets/sec. For the token ring LAN, this threshold value is evaluated to be much lower; at about 700 packets/sec. Moreover, compared to Ethernet, higher packet delays are encountered for lower traffic loads.

This is due to the fact that as the traffic on the ring is increased, many nodes will have frames to send that will be delayed until the token is passed around the ring. Each node will in this fashion conserve the token until it transmits all its frames and consequently queues will form in all other nodes increasing as such their delay.

Part 2: Routing Experiments

Before running any experiments on how OSPF protocol performs while routing traffic across the network, we examined the time it takes for the protocol to converge – establish a routing table that the network will follow. In figure 2, we provide the obtained results. As it can be seen from that figure, the traffic sent becomes periodic after 50s. In the period from 0 to 50s the OSPF protocol sends messages across the network, and by employing a specific criteria decides which path traffic will follow. The criteria that was set by default is based on the bandwidth of the links. The cost of the links are inversely proportional to the bandwidth. In a such setup the routing table looked as described in Table 3.

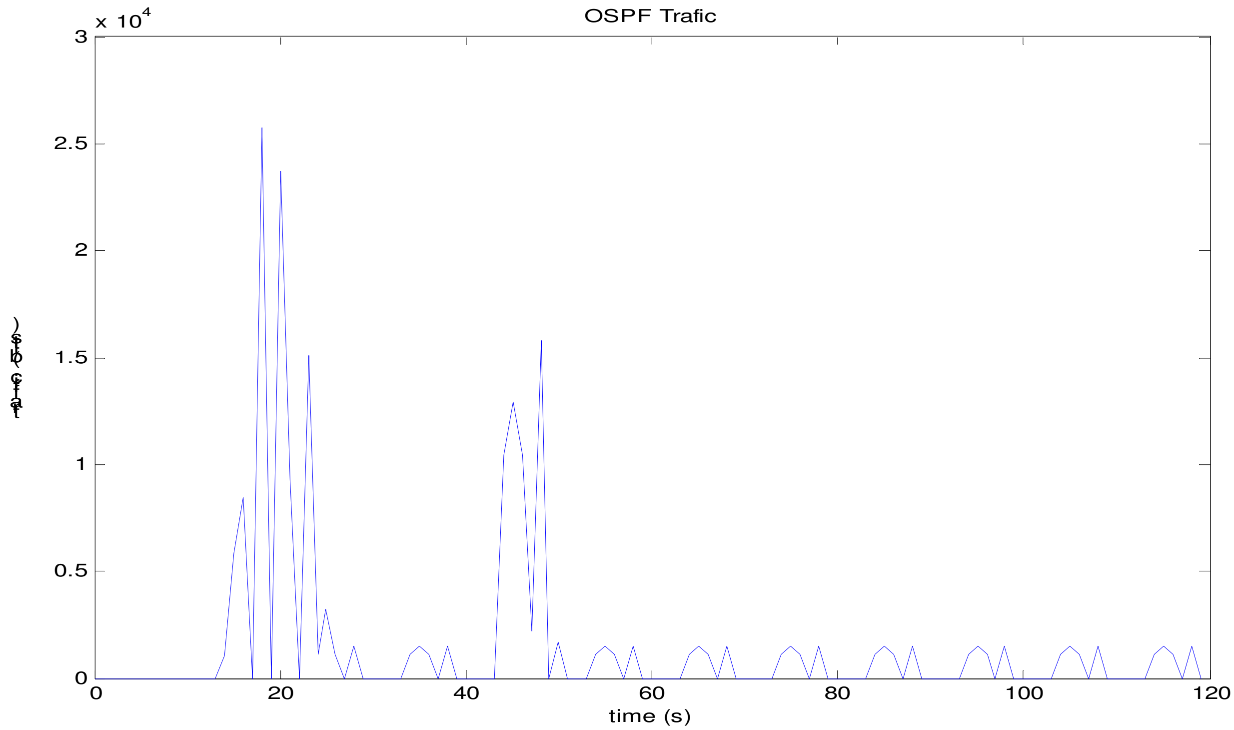

Figure 6: Traffic sent by the OSPF protocol and the time it takes to converge.

The OSPF is a publicly available routing protocol used for intra-AS routing[1]. It is a link state protocol that uses flooding of link state information and Dijkstra least-cost path algorithm. OSPF does not set the link costs explicitly that are used locally at the router to run Dijkstra algorithm. Instead link costs are assigned by the network administrator. Setting the link costs to be inversely proportional to link capacity is used to discourage traffic from using low-bandwidth links. This is not an optimum solution, as it will be seen from the sections to follow. In Figure 6 we still observe some activity of the protocol even after 50s. This is due to the fact that the protocol periodically broadcasts link's state even if the link's state hasn't changed. The information that OSPF sends across the network are contained in the OSPF messages and are carried by the IP. Due to this OSPF must implement its own mechanism to achieve reliable transfers.

Source	Destination	Path Taken
A	C	A-C
A	D	A-D

Source	Destination	Path Taken
A	C	A-B-C
A	D	A-B-C-D

B	A	B-A
C	A	C-A
D	B	D-C-B
D	C	D-C

Table 3. Routing table given the initial conditions

B	A	B-A
C	A	C-B-A
D	B	D-C-B
D	C	D-C

Table 4. Routing table after changing the properties of the physical link between A and B from 0C1 to 0C3

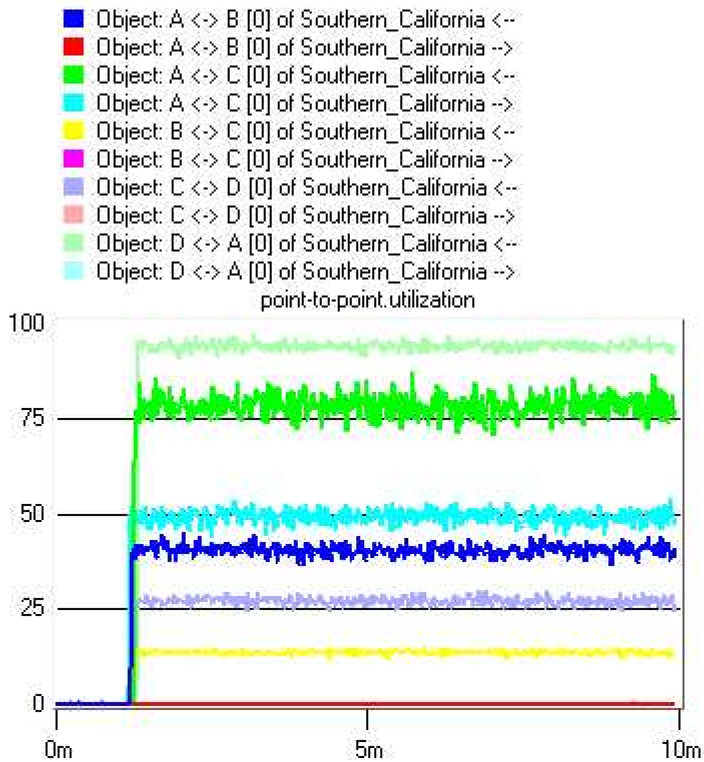

Figure 7. Link utilization from point to point given the initial conditions

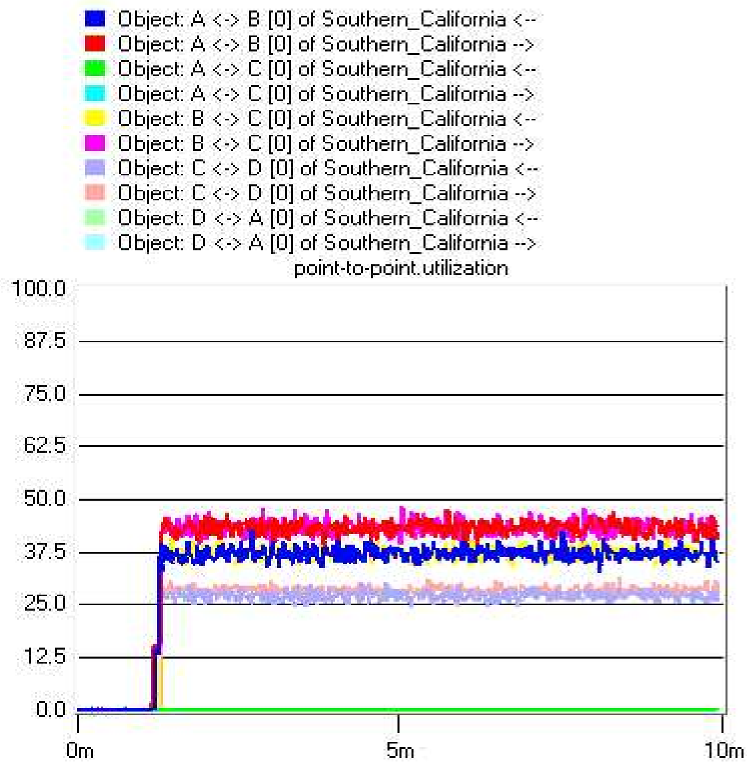

Figure8. Link utilization from point to point after changing the A-B link from OC1 to OC3

From Figure 7, we observe that with initial (default) settings utilization of link A-D reaches almost 100%. This result is consistent since we know that A-D is a DS-3 type link and routing of 42Mbps is done across it. It can be concluded that this link represents the most heavily used link in the network.
By upgrading the capacity of one of the links we have tried to reduce the utilization. One level upgrade was available to us. By examining the links, we concluded that upgrading from DS-3 to OC-1 will not be sufficient because the effective bandwidth will go up from 44.736 Mbps to 49.36 Mbps which is only 10% increase. Instead, by upgrading the OC -1 link to OC – 3 will result in more than 200% improvement. This improvement will not directly reflect on the total link utilization, but will certainly reduce it substantially. Now, since the OSPF protocol was set to weight the links inversely proportional to their capacity, we have obtained a new routing table that is presented in Table 4. As we can see no traffic is routed across the DS-3 type links, instead all of the traffic flows across the OC-3. The link utilization has decreased by almost 50%, which is shown in Figure 8.

In addition to changing the properties of the physical link, we have tried to do even better by manually assigning link costs at the router interfaces. Unfortunately, this could not be achieved due to the fact that DS-3 links will have to be utilized. From Table 1, we see that the lowest traffic in the network is 20 Mbps. If we are to transfer this much data across a DS-3 link we would already be at 45%. Then the 42Mbps going from router A to D will have to get routed across the A-B OC3 link which will bring the utilization up to 86% which is certainly much worse comparing to the scenario where we upgraded the physical link.

In the following part we have examined a slightly different situation. A failure of a link and how do the traffic demands change. The data in Table 4. was collected from the IP protocol demands. We see that no traffic is now routed across the B-C link since it practically does not exist anymore. Instead the only traffic that was going across that link – D-B is now being routed across A-B link. This brings network to its maximum utilization level, which is certainly not a desirable condition. The quality of service will be tremendously decreased.

Source	Destination	Path Taken
A	C	A-C
A	D	A-D
B	A	B-A
C	A	C-A
D	B	D-A-B
D	C	D-C

Table 4. Routing table after link B-C fails with initial conditions (A-B link is OC 1)

Figure 9, shows what is the condition of the network provided in Scenario 2 after B-C link fails. We observe that diagonal D-B link is being used at 100%, and as such requires manual link cost assignment. Currently the data is routed across D-B and a C-D link. In our assessment of the network, we have concluded that data from router C should be routed via router A instead of router D. To achieve this setup, we have assigned high link cost at the transfer port at router C that is linked to router D. Results are presented in Figure 10. Utilization of the network has gone down to 50%. As such, we can say that in case of the failure the network will not be overloaded, and it will still be providing quality of service.

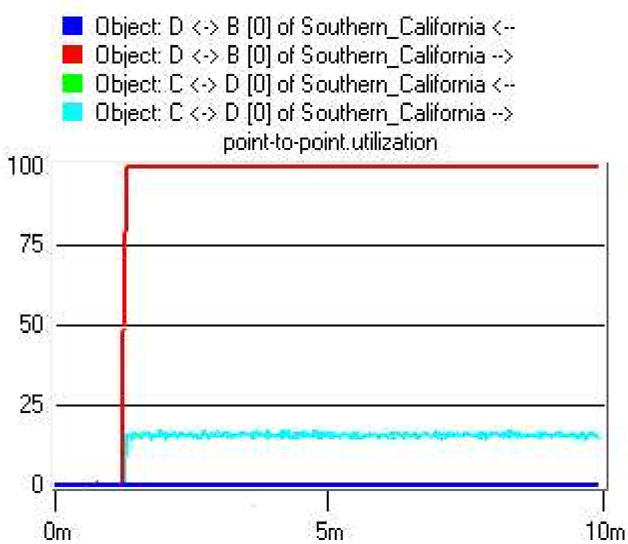

Figure 9. Link point to point utilization once the B-C link fails in Scenario 2

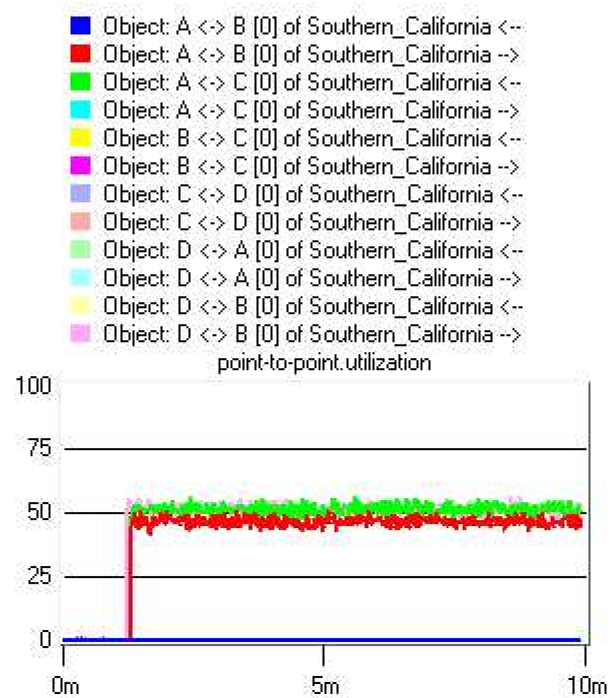

Figure 10. Link point to point utilization after we have manually assigned link costs in Scenario 2.

Conclusion

This experiment permitted us to explore multiple access control protocols in the context of LANs on one hand and the detailed operation of the OSPF routing protocol on the other. We observed how propagation delay affects the effectiveness of an Ethernet as well as the token ring LAN and furthermore how packet delay was increased as more traffic became present on the networks. These experiments drove us to the conclusion that a much higher effectiveness as well as shorter delays are experienced on a Ethernet network compared to its token ring counterpart. Elsewhere, we observed the dynamics of the intra-AS routing process in the OSPF protocol. We observed how attributing link costs according to bandwidth results in an overloading of the network with the occurrence of a link failure. We successively experimented how user controlled link cost assignment can resolve this latter quandary.

References

1. James F. Kurose, Keith W. Ross, *Computer Networking: A Top-Down Approach Featuring the Internet,* Addison Wesley, 2nd Edition, 2002
2. Alberto Leon-Garcia and Indra Widjaja, *Communication Networks,* McGraw- Hill, first ed., 2000.
3. D. Bertsekas and R. Gallagher, *Data Networks, 2nd Ed.,* Prentice Hall, Englewood Cliffs, NJ, 1991.